arXiv:0902.1661v1 [cs.CC] 10 Feb 2009

Even Faster Exact Bandwidth *

Marek Cygan and Marcin Pilipczuk
{cygan,malcin}@mimuw.edu.pl

Department of Mathematics, Computer Science and Mechanics,
University of Warsaw, Warsaw, Poland

Abstract

We deal with exact algorithms for Bandwidth, a long studied NP-hard problem. For a long time nothing better than the trivial $O^*(n!)$[1] exhaustive search was known. In 2000, Feige an Kilian [5] came up with a $O^*(10^n)$-time algorithm. Recently we presented algorithm that runs in $O^*(5^n)$ time and $O^*(2^n)$ space.

In this paper we present a major modification to our algorithm which makes it run in $O(4.83^n)$ time with the cost of $O^*(4^n)$ space complexity. This modification allowed us to perform Measure & Conquer analysis for the time complexity which was not used for such types of problems before.

1 Introduction

In this paper we focus on exact exponential-time algorithms for the Bandwidth problem. Let $G = (V, E)$ be an undirected graph, where $n = |V|$ and $m = |E|$. For a given one-to-one function $\pi : V \to \{1, 2 \ldots, n\}$ (called *ordering*) its *bandwidth* is the maximum difference between positions of adjacent vertices, i.e. $\max_{uv \in E} |\pi(u) - \pi(v)|$. The *bandwidth* of the graph, denoted by $\mathrm{bw}(G)$, is the minimum bandwidth over all orderings. The Bandwidth problem asks to find an ordering with bandwidth $\mathrm{bw}(G)$. One can consider a decision version of the Bandwidth problem. Precisely, we assume that the input for our problem contains additionally an integer b, $1 \leq b < n$. An ordering of V with bandwidth at most b is called a *b-ordering*. In the decision version we ask if there exists a b-ordering and if that is the case, finding it.

A very short summary of what is known about the Bandwidth problem follows. On special families of graphs $\mathrm{bw}(G)$ can be computed in polynomial time [1, 8]. However, in general Bandwidth is NP-hard even on some subfamilies of trees [7, 9]. Moreover Unger [10] showed that Bandwidth problem does not belong to APX even in a very restricted case when G is a caterpillar, i.e. a very simple tree. It is also hard for any

*This research is partially supported by a grant from the Polish Ministry of Science and Higher Education, project N206 005 32/0807.

[1] By O^* we denote standard big O notation but omitting polynomial factors.

fixed level of the W hierarchy [2]. The best known polynomial-time approximation, due to Feige [4], has $O(\log^3 n\sqrt{\log n \log\log n})$ approximation guarantee. On WG 2008 we presented [3] exact algorithm that runs in $O^*(5^n)$ time and $O^*(2^n)$ space.

In this paper we enhance the previous algorithm to make it run in $O(4.83^n)$ time. However, the cost of this change is $O^*(4^n)$ space complexity, which makes this result purely theoretical. In Section 2 we describe the enhanced algorithm. In Section 3 we do Measure & Conquer analysis (method introduced by Fomin et al. in [6]) to obtain $O(4.83^n)$ time bound. We find this analysis interesting, because Measure & Conquer method at the first glance does not fit for the BANDWIDTH problem at all.

2 The algorithm

The algorithm consists of two phases. First, we generate partial assignments of vertices to positions: we do not assign precise position to a vertex, but a segment of length $2(b+1)$ or $4(b+1)$ of possible positions. We do this in every possible way. In the second phase for every generated segment assignment we check whether there exists a precise assignment of vertices to positions (i.e., an ordering), consistent with the partial assignment.

From this point we assume, that the graph G has at least two vertices and it is connected (if G is not connected we may find b-orderings of each connected component of G in an independent manner). Let us choose any (but fixed for the whole algorithm) spanning tree $\mathcal{D}$ of the graph G.

2.1 First phase: generating segment assignments

Definition 2.1. *A* segment *is a nonempty set of consecutive positions which has form of* $\{i(b+1)+1, i(b+1)+2, \ldots, j(b+1)\} \cap \{1, 2, \ldots, n\}$ *for some integers* $i < j$*. We say that this segment has index* (i, j) *and denote it as* $\Theta_{(i,j)}$*. For the sake of simplicity we define* $\Theta_i = \Theta_{(i,i+1)}$ *and call such segments* base segments.

Definition 2.2. *A* segment assignment *is a function* ϕ *assigning a segment to every vertex such that the following conditions hold:*

1. *Every leaf of the spanning tree* $\mathcal{D}$ *is assigned to a segment of size* $4(b+1)$*, i.e., segment* $\Theta_{(i,i+4)}$ *for some integer* i*.*
2. *Every inner vertex of* $\mathcal{D}$ *is assigned to a segment of size* $2(b+1)$*, i.e., segment* $\Theta_{(i,i+2)}$ *for some integer* i*.*
3. *For every edge* uv *in* $\mathcal{D}$*, if vertex* u *is the parent of the vertex* v *and* v *is an inner vertex, where* $\phi(u) = \Theta_{(i,i+2)}$ *and* $\phi(v) = \Theta_{(j,j+2)}$*, then* i *and* j *differ by exactly one.*
4. *For every leaf* v*, if* u *is the parent of* v *in* $\mathcal{D}$ *and* $\phi(v) = \Theta_{(i,i+2)}$*, then* $\phi(u) = \Theta_{(i-1,i+3)}$*.*

We say that a segment assignment is consistent *with an ordering π if for every vertex v position $\pi(v)$ belongs to the segment $\phi(v)$.*

Lemma 2.3. *Let π be a b-ordering. In any segment assignment ϕ consistent with the ordering π, for every edge uv, if $\phi(u) = \Theta_{(i,j)}$ and $\phi(v) = \Theta_{(k,l)}$ then $j \geq k$ and $l \geq i$.*

Proof. If $j < k$, then there is a gap of size at least $b + 1$ between $\phi(u)$ and $\phi(v)$. Since π is consistent with the segment assignment, edge uv is longer than b, contradiction. Similar argument proves that $l \geq i$. □

Lemma 2.4. *Let π be a b-ordering. There exists a segment assignment consistent with π.*

Proof. As a proof we present simple construction of the segment assignment. Let u_0 be the root of $\mathcal{D}$. Lets assign it to any segment of length $2(b + 1)$ containing $\pi(u_0)$ (there are exactly two possible segments). Then assign segments to vertices in the root-to-leaf order. Let u be an unassigned vertex with parent v and let $T_{(i,i+2)}$ be the segment assigned to v. Since $i(b + 1) + 1 \leq \pi(v) \leq (i + 2)(b + 1)$ and π is a b-ordering,

$$\pi(u) \in \Theta_{(i-1,i+3)} = \Theta_{(i-1,i+1)} \cup \Theta_{(i+1,i+3)}.$$

We can assign v to $\Theta_{(i-1,i+3)}$ or to one of the segments $\Theta_{(i-1,i+1)}$ and $\Theta_{(i+1,i+3)}$, containing $\pi(u)$, depending whether v is an inner vertex or a leaf. □

Our goal in the first phase is to generate a set of segment assignments such that for every b-ordering there exists generated segment assignment consistent with it. In the second phase we check for every segment assignment whether consistent b-ordering exists. As a result we check if there exists any b-ordering of the given graph G.

The first phase is as follows:

1. Assign root u_0 of the tree $\mathcal{D}$ to any valid segment of size $2(b + 1)$.
2. Recursively assign other vertices in the root-to-leaf order in tree $\mathcal{D}$. Given vertex u with parent v assigned to segment $\Theta_{(i,i+2)}$ assign u to one of the segments $\Theta_{(i-1,i+1)}$, $\Theta_{(i+1,i+3)}$ if u is an inner vertex, or to $\Theta_{(i-1,i+3)}$ when u is a leaf (see Figure 1).
3. For every generated assignment check condition from Lemma 2.3: in other words accept assignment iff for every edge uv in G segments assigned to vertices u and v are not too far from each other.

Note that Lemma 2.4 implies that for every b-ordering π there is a generated segment assignment consistent with π. Note that this is exactly the assignment described in the proof of Lemma 2.4.

The second phase is fed with generated segment assignments on-line. Therefore the first phase uses polynomial space. In Step 2 of the algorithm we have two possibilities for every inner non-root vertex. For the root we have $\lceil \frac{n}{b+1} \rceil + 1 \leq n + 1$ possibilities

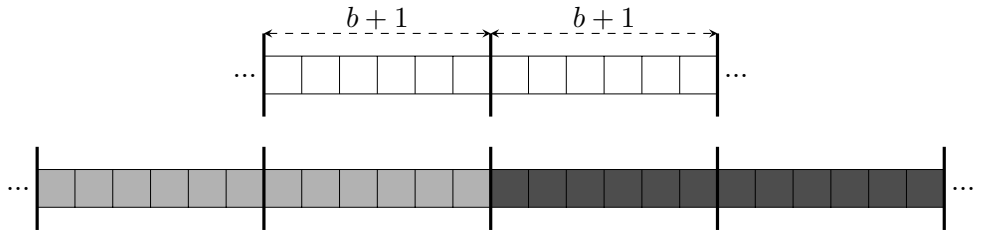


Figure 1: Upper part of the picture shows the segment $\Theta_{i,i+2}$ assigned to the vertex v. Lower part contains possible positions for the vertex u (child of v in the tree D) covered by two segments $\Theta_{i-1,i+1}$ and $\Theta_{i+1,i+3}$.

(additional 1 comes from segment $\Theta_{-1,1}$, which is also correct segment according to the Definition 2.1). We do not have a choice for leaves, thus the algorithm generates at most $(n+1)2^{n-1}$ assignments. This algorithm uses polynomial time for every generated assignment, so this phase works in $O^*(2^n)$ time.

2.2 Second phase: depth-first search

In this phase for every generated segment assignment we check if there exists a b-ordering consistent with the segment assignment. Let us denote this segment assignment by ϕ. This phase is very similar to the second phase of our $O^*(5^n)$ algorithm [3], but adapted to the segments of size $2(b+1)$ and $4(b+1)$. This difference allows us to use Measure & Conquer method in complexity analysis. First, we recall some key definitions and facts from the $O^*(5^n)$ algorithm.

We assign a vertex to each position one by one, but the main idea is the order in which we fill those positions in. For every position i, let $\texttt{segment}(i) = \lceil \frac{i}{b+1} \rceil$ be the base segment number of this position, and let $\texttt{color}(i) = ((i-1) \mod (b+1)) + 1$ be the index of the position in its base segment, which we call the *color* of this position. Note that the color of position is the remainder of this number modulo $b+1$, but in the range $[1, b+1]$ instead of $[0, b]$.

Let us sort positions lexicographically according to pairs $(\texttt{color}(i), \texttt{segment}(i))$. To each of those positions we assign a vertex, in exactly this order. We call this ordering *the color order* of positions.

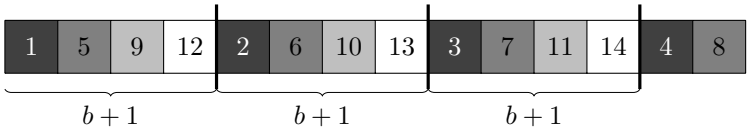


Figure 2: Color order of positions for $n = 14$ and $b = 3$.

Following lemma is the key observation in our algorithm.

Lemma 2.5. *Ordering π is a b-ordering iff for every edge uv such that* $\texttt{segment}(\pi(u)) < \texttt{segment}(\pi(v))$ *we have* $\texttt{segment}(\pi(u))+1 = \texttt{segment}(\pi(v))$ *and* $\texttt{color}(\pi(u)) > \texttt{color}(\pi(v))$.

Proof. Since π is a b-ordering, for every edge uv we have $|\texttt{segment}(\pi(u))-\texttt{segment}(\pi(v))| \leq 1$. If $\texttt{segment}(\pi(u)) = \texttt{segment}(\pi(v))$ then uv is not longer than b. Otherwise, suppose w.l.o.g. that $\texttt{segment}(\pi(u)) + 1 = \texttt{segment}(\pi(v))$. Note that the distance between positions with the same color in the neighboring segments is exactly $b + 1$, so uv is not longer than b iff u has greater color than v (see Figure 3). □

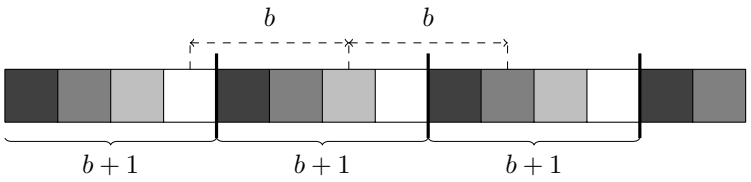


Figure 3: Picturable proof of the Lemma 2.5

Corollary 2.6. *Ordering π, is a b-ordering iff for every edge uv satisfying* $\texttt{segment}(\pi(u))+1 = \texttt{segment}(\pi(v))$ *vertex u is assigned to greater position in the color order than vertex v and for every edge uv values* $\texttt{segment}(\pi(u))$ *and* $\texttt{segment}(\pi(v))$ *differ by at most one.*

Definition 2.7. *By* state candidate *we denote partial function $s : V \to \{\Theta_i : 0 \leq i < \lceil \frac{n}{b+1} \rceil\}$, i.e., assigning base segments to some of the vertices, such that for every vertex v if $s(v)$ is defined then $s(v) \subset \phi(v)$ (where ϕ is the fixed assignment for this phase).*

Note, that for an inner vertex v of the tree $\mathcal{D}$ there are two possible values for the state candidate s and there are four possible values for v being a leaf.

Definition 2.8. *A* state *is a state candidate s satisfying the following conditions:*

1. *Vertices from* $\mathrm{dom}(s)$ *can be assigned to the first* $|\mathrm{dom}(s)|$ *positions in the color order, in a way compatible with s.*
2. *For every edge uv one of the following holds.*
 (a) *Both $s(u)$ and $s(v)$ are undefined.*
 (b) *Exactly one value among $\{s(u), s(v)\}$ is defined; say $s(v) = \Theta_i$ and $s(u)$ is undefined. Then if $\phi(u) = \Theta_{(k,l)}$ we have $k \leq i$.*
 (c) *Both $s(u)$ and $s(v)$ are defined; say $s(v) = \Theta_i$ and $s(v) = \Theta_k$. Then $|i-k| \leq 1$.*

Definition 2.9. *We say that a state s' is an* extension *of some state s if there exists vertex v such that:*

1. *$s(v)$ is undefined and $s'(v)$ is defined;*

2. $\mathrm{dom}(s') = \mathrm{dom}(s) \cup \{v\}$ *and* $s = s'|_{\mathrm{dom}(s)}$*;*

3. *If* uv *is an edge in the* **whole graph** G *and* $s'(u) = \Theta_k$ *and* $s'(v) = \Theta_i$*, then* $k - 1 \leq i \leq k$*.*

Lemma 2.10. *Let* π *be a* b*-ordering compatible with the given segment assignment* ϕ*. By* s_k *for* $0 \leq k \leq n$ *we denote the state candidate, assigning* $\texttt{segment}(\pi(v))$ *to every vertex* v *assigned to one of the first* k *positions in the color order (i.e.,* $|\mathrm{dom}(s_k)| = k$*). Then every* s_k *is a state and for every* $0 \leq k < n$ *state* s_{k+1} *is an extension of state* s_k*.*

Proof. This is a straightforward corollary from Corollary 2.6. Point 1 of the state definition is obviously satisfied by the definition of s_k. Now lets look at any s_k and edge uv. If both $s_k(u)$ and $s_k(v)$ are undefined, Point 2 of the state definition is satisfied. If both are defined, Corollary 2.6 clearly implies Point 2. Assume then that $s_k(v) = \Theta_i$ and $s_k(u)$ is undefined. By construction of state candidates s_k, $\pi(u)$ is later in the color order than $\pi(v)$. Therefore, by Corollary 2.6, $\texttt{segment}(\pi(u)) \leq \texttt{segment}(\pi(v))$. Since $\pi(u) \in \texttt{segment}(\pi(u)) \subset \phi(u)$, it implies the condition in Point 2 of the state definition.

Let us now prove that s_{k+1} is an extension of the state s_k. Points 1 and 2 of the extension definition follow directly from the construction of s_k. Note that if v is the vertex defined by s_{k+1} but not by s_k, then $\pi(v)$ is later in the color order than any u defined by s_k. Therefore, if we take any uv as in Point 3 of the extension definition, by Corollary 2.6 the condition in Point 3 is satisfied. □

Lemma 2.11. *Let* s_k *for* $0 \leq k \leq n$ *be a set of states such that* s_{k+1} *is an extension of state* s_k *for every* $0 \leq k < n$*. Let* v_k *be the vertex defined by* s_k*, but not by* s_{k-1}*. Then ordering* π *which assigns position* k *in the color order to the vertex* v_k *is a* b*-ordering.*

Proof. We use Corollary 2.6 once again. Let us prove the thesis by contradiction. Assume that edge uv is longer than b and w.l.o.g. $\pi(u) < \pi(v)$. Since $\pi(u) + b < \pi(v)$, $\texttt{segment}(\pi(u)) < \texttt{segment}(\pi(v))$. If $\texttt{segment}(\pi(u)) + 1 < \texttt{segment}(\pi(v))$, then s_n is not a state due to Point 2c of the state definition. Therefore $\texttt{segment}(\pi(u)) + 1 = \texttt{segment}(\pi(v))$ and, by Corollary 2.6, $\pi(u)$ is earlier in the color order than $\pi(v)$. But then there exists s_k for which $s_k(u)$ is defined, $s_k(v)$ is undefined but $s_{k+1}(v)$ is defined. Thus we have contradiction with Point 3 of the extension definition for the edge uv. □

The algorithm for second phase is now quite clear. Given segment assignment ϕ, we do depth-first search over states, looking for sequence of states and extensions as in Lemmas 2.10 and 2.11. More precisely:

1. Start with a state $\emptyset$.

2. Having state s with $|\mathrm{dom}(s)| = k < n$, try to create extension s' of s in every possible way, i.e., try to assign every undefined vertex in s to the base segment, where $(k+1)$-th position in the color order lies.

3. If we reach state s with $|\mathrm{dom}(s)| = n$, construct the b-ordering π using the DFS stack (that is, states $\emptyset = s_0, s_1, \ldots, s_n = s$) and return it.

4. If no state s with $|\mathrm{dom}(s)| = n$ is reached, return that there is no b-ordering compatible with ϕ.

Note that Lemmas 2.10 and 2.11 imply that this algorithm returns correct b-ordering in Step 3 and there are no correct b-orderings if the algorithm reaches Step 4. Therefore we proved that this algorithm is correct. In Section 3 we prove the $O(4.83^n)$ time bound and $O^*(4^n)$ space bound.

3 Time and space analysis

3.1 Memory bound

This analysis is fairly easy. Note that only non-polynomial space used in the algorithm is the space used to track visited states in the depth-first search. We try to bound the number of visited states for one fixed segment assignment ϕ by 4^n. This implies $O^*(4^n)$ space bound.

Lemma 3.1. *In one run of the second phase, the algorithm visits at most $3^{n-L}4^L \leq 4^n$ states, where L is the number of leaves in the tree $\mathcal{D}$.*

Proof. Let s be a visited state and v be a vertex.

If v is an inner vertex of $\mathcal{D}$ with $\phi(v) = \Theta_{(i,i+2)}$, then $s(v)$ is either undefined or $s(v) = \Theta_i$ or $s(v) = \Theta_{i+1}$ - three possibilities.

If v is a leaf of $\mathcal{D}$ with $\phi(v) = \Theta_{(i-1,i+3)}$, then $s(v)$ is either undefined or one of the base segments Θ_{i-1}, Θ_i, Θ_{i+1}, Θ_{i+2} - five possibilities, too many. However, lets look at the parent u of the vertex v in the tree $\mathcal{D}$. By construction, $\phi(u) = \Theta_{(i,i+2)}$. Note that, by the definition of the state, since uv is an edge of G:

1. If $s(u)$ is undefined, then $s(v)$ cannot be Θ_{i-1}.
2. If $s(u) = \Theta_i$, then $s(v)$ cannot be Θ_{i+2}.
3. If $s(u) = \Theta_{i+1}$, then $s(v)$ cannot be Θ_{i-1}.

Therefore, in every case, there are only four possibilities for a leaf v. This leads to $3^{n-L}4^L$ bound for the number of valid states. □

3.2 Time bound

Note that search for possible expansions of given state and checking if a state candidate is a state can be done in polynomial time. Therefore the time used by the whole algorithm is bounded by $O((s(G) + 2^n)n^\gamma)$ where γ is a constant and $s(G)$ is the number of visited states. The 2^n factor is due to the first phase of the algorithm. Now we focus on bounding $s(G)$, i.e., the number of visited states in the run of the whole algorithm.

Lemma 3.2. *The algorithm in the whole run visits at most $3(n+1)\kappa^n$ states for some constant $\kappa < 4.83$.*

Proof. Let us do indepth analysis of the number of possible states similar, but more broad, to that in the proof of Lemma 3.1. Let s be a state visited while considering segment assignment ϕ.

Let u_0 be the root of $\mathcal{D}$. Then $\phi(u_0)$ has at most $\lceil \frac{n}{b+1} \rceil + 1$ possible values. If $\phi(u_0) = \Theta_{(i,i+2)}$, then $s(u_0)$ is either undefined or equal to Θ_i or Θ_{i+1}. In total, $3(n+1)$ possibilities for u_0.

Let u be a parent of a leaf v. Then there are at most four possibilities to choose for $\phi(v)$ and $s(v)$, knowing $\phi(u)$ and $s(u)$. The analysis is the same as in the proof of Lemma 3.1.

Let now v be an inner vertex with parent u. Let $\phi(u) = \Theta_{(i,i+2)}$. Then $\phi(v)$ is either $\Theta_{(i-1,i+1)}$ or $\Theta_{(i+1,i+3)}$, by the way we construct ϕ. The following restrictions are implied by the fact that uv is an edge of G and by the state definition.

A. If $s(u)$ is undefined:

 (a) If $\phi(v) = \Theta_{(i-1,i+1)}$, then $s(v)$ is undefined or $s(v) = \Theta_i$. State definition forbids possibility that $s(v) = \Theta_{i-1}$.

 (b) If $\phi(v) = \Theta_{(i+1,i+3)}$, then $s(v)$ is undefined, $s(v) = \Theta_{i+1}$ or $s(v) = \Theta_{i+2}$.

 In total, **five** possibilities for $\phi(v)$ and $s(v)$.

B. If $s(u) = \Theta_i$:

 (a) If $\phi(v) = \Theta_{(i-1,i+1)}$, then $s(v)$ is undefined, $s(v) = \Theta_{i-1}$ or $s(v) = \Theta_i$.

 (b) If $\phi(v) = \Theta_{(i+1,i+3)}$, then $s(v) = \Theta_{i+1}$. State definitions forbids possibilities that $s(v)$ is undefined and that $s(v) = \Theta_{i+2}$.

 In total, **four** possibilities for $\phi(v)$ and $s(v)$.

C. If $s(u) = \Theta_{i+1}$:

 (a) If $\phi(v) = \Theta_{(i-1,i+1)}$, then $s(v)$ is undefined or $s(v) = \Theta_i$. State definition forbids possibility that $s(v) = \Theta_{i-1}$.

 (b) If $\phi(v) = \Theta_{(i+1,i+3)}$, then $s(v)$ is undefined, $s(v) = \Theta_{i+1}$ or $s(v) = \Theta_{i+2}$.

 In total, **five** possibilities for $\phi(v)$ and $s(v)$.

In every option we got at most five options for $\phi(v)$ and $s(v)$ values for every inner, non-root vertex v. Together with four possibilities for leaves and $3(n+1)$ for root, this proves that the algorithm visits at most $3(n+1)5^n$ states.

However, there are few places where we have four, not five possibilities for $\phi(v)$ and $s(v)$: when v is a leaf or when the parent of u is assigned by the state to the left (smaller position numbers) half of its segment (Option B). Moreover, in every moment, when we have five possibilities for $\phi(v)$ and $s(v)$, vertex v might be assigned to the left (smaller position numbers) half of its segment, which gives us Option B for the analysis of children

of v. This leads us to the conclusion that we can use Measure & Conquer technique to obtain better bound.

The Measure & Conquer method was introduced by Fomin, Grandoni and Kratsch (see [6]). As in the above analysis, we analyze vertices in the root-to-leaf order. We use Measure & Conquer method to estimate number of possible states.

Let us consider hypothetical state generator that generates possible states (i.e., pairs of functions ϕ and s), by analyzing the tree $\mathcal{D}$ in the root-to-leaf order. The generator first sets $\phi(u_0)$ and $s(u_0)$ in every possible way (at most $3(n+1)$ ways). Then, while analyzing vertex v with already set $\phi(u)$ and $s(u)$ for parent u of v, it assigns $\phi(v)$ and $s(v)$ in every possible way, keeping in mind limitations described above, both for v being a leaf and inner vertex. We use Measure & Conquer to estimate number of generated states by this generator.

At any step of the generator we measure the weight of the current problem instance, i.e., already constructed functions ϕ and s. The weight of the instance is the sum of weights of vertices. Let $\alpha, \beta \in [0,1]$ be constants to be defined lated. The weight of the vertex v is:

I. 0, if v is already analyzed;

II. 1, if v is not analyzed and parent of v is not analyzed or v is the root of $\mathcal{D}$;

III. 1, if v is not analyzed, parent u of v is analyzed and $s(u)$ is undefined;

IV. α, if v is not analyzed, parent u of v is analyzed and $\phi(u) = \Theta_{(i,i+2)}$ and $s(u) = \Theta_i$ for some integer i.

V. β, if v is not analyzed, parent u of v is analyzed and $\phi(u) = \Theta_{(i,i+2)}$ and $s(u) = \Theta_{i+1}$ for some integer i.

Let $T(w)$ be a bound for number of states generated by the generator from point, where the size of the instance is at most w. Now we estimate T, using aforementioned limitations for generator choices.

Let v be a non-root vertex currently analyzed by the generator, where vertex u is its parent. If v is a leaf, there are always four possibilities for $\phi(v)$ and $s(v)$, independent of whether v falls into Category III, IV, V, i.e., whether v weights 1, α or β. Therefore

$$T(w) \leq \max(4T(w-1), 4T(w-\alpha), 4T(w-\beta)).$$

Now lets look at the case when v is an inner vertex. Let $\phi(u) = \Theta_{(i,i+2)}$.

1. If $s(u)$ is undefined, v has got weight 1 and we have five possibilities for $\phi(v)$ and $s(v)$. In two of them, v falls under Category III for children of v, in two — under Category V, and in one — under Category IV. Vertex v becomes analyzed and has weight 0. Since there is at least one child of vertex v, the following bound holds in this case:

$$T(w) \leq 2T(w-1) + 2T(w-1-(1-\beta)) + T(w-1-(1-\alpha)).$$

2. If $s(u) = \Theta_i$, v has got weight α and we have four possibilities for $\phi(v)$ and $s(v)$. In one of them, v falls under Category III for children of v, in one — under Category V, and in two — under Category IV. Vertex v becomes analyzed and has weight 0. Since there is at least one child of vertex v, the following bound holds in this case:

$$T(w) \leq T(w-\alpha) + T(w-\alpha-(1-\beta)) + 2T(w-\alpha-(1-\alpha)).$$

3. If $s(u) = \Theta_{i+1}$, v has got weight β and we have five possibilities for $\phi(v)$ and $s(v)$. In two of them, v falls under Category III for children of v, in two — under Category V, and in one — under Category IV. Vertex v becomes analyzed and has weight 0. Since there is at least one child of vertex v, the following bound holds in this case:

$$T(w) \leq 2T(w-\beta) + 2T(w-\beta-(1-\beta)) + T(w-\beta-(1-\alpha)).$$

By searching the space of possible values α and β and by solving the above equations numerically, we got that for $\alpha = 0.8805$ and $\beta = 1$ function $T(w)$ is bounded by κ^n for $\kappa \sim 4.828485 < 4.83$. This completes the proof. □

The following theorem is a straightforward corollary from Lemmas 3.1 and 3.2.

Theorem 3.3. *There exists an algorithm that solves the* BANDWIDTH *problem in* $O(4.83^n)$ *time and* $O^*(4^n)$ *space.*

References

[1] S. Assman, G. Peck, M. Syslo, and J. Zak. The bandwidth of caterpillars with hairs of length 1 and 2. *SIAM J. Algebraic Discrete Methods*, 2:387–393, 1981.

[2] H. L. Bodlaender, M. R. Fellows, and M. T. Hallett. Beyond NP-completeness for problems of bounded width: Hardness for the w hierarchy (extended abstract). *ACM Symposium on Theory of Computing*, pages 449–458, 1994.

[3] M. Cygan and M. Pilipczuk. Faster exact bandwidth. *WG 2008 Proceedings*, 2008.

[4] U. Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. Syst. Sci.*, 60(3):510–539, 2000.

[5] U. Feige. Coping with the NP-hardness of the graph bandwidth problem. *SWAT*, pages 10–19, 2000.

[6] Fedor V. Fomin, Fabrizio Grandoni, and Dieter Kratsch. Measure and conquer: a simple $O(2^{0.288n})$ independent set algorithm. In *SODA'06*, pages 18–25, 2006.

[7] M. Garey, R. Graham, D. Johnson, and D. Knuth. Complexity results for bandwidth minimization. *SIAM J. Appl. Math.*, 34:477–495, 1978.

[8] D. Kleitman and R. Vohra. Computing the bandwidth of interval graphs. *SIAM J. Discrete Math.*, 3:373–375, 1990.

[9] B. Monien. The bandwidth-minimization problem for caterpillars with hairlength 3 is np-complete. *SIAM J. Algebraic Discrete Methods*, 7:505–512, 1986.

[10] W. Unger. The complexity of the approximation of the bandwidth problem. *FOCS*, pages 82–91, 1998.